

AD-A126 340 OBJECT DETECTION BY TWO-DIMENSIONAL LINEAR PREDICTION
(U) MASSACHUSETTS INST OF TECH LEXINGTON LINCOLN LAB

118

OBJECT DETECTION BY TWO-DIMENSIONAL LINEAR PREDICTION
(U) MASSACHUSETTS INST OF TECH LEXINGTON LINCOLN LAB


T F QUATIERI 26 JAN 83 TR-632 ESD-TR-82-109

UNCLASSIFIED


F 19628-80-C-0002

F/G 12/1

NL



END
4-83



MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

AD A 126340

(12)

The work reported in this document was performed at Lincoln Laboratory, a center for research sponsored by Massachusetts Institute of Technology, with the support of the Department of the Air Force under Contract D11000-60-C-0001.

This report may be reproduced as easily made of U.S. Government agencies.

The views and conclusions contained in this document are those of the author(s) and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the United States Government.

The Public Affairs Office has reviewed this report, and it is available to the National Technical Information Service, where it will be available to the general public, including foreign nations.

JL-9

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
LINCOLN LABORATORY

**OBJECT DETECTION
BY TWO-DIMENSIONAL LINEAR PREDICTION**

T.F. QUATIERI
Group 27

TECHNICAL REPORT 632

28 JANUARY 1963

Approved for public release; distribution unlimited.

LEXINGTON

MASSACHUSETTS

The originals of the color figures appearing in this document are color photographic prints. Recipients of this document desiring a similar quality reproduction may obtain the full set at the cost of \$20.00, by submitting a check payable to Massachusetts Institute of Technology to:

Lincoln Laboratory
Massachusetts Institute of Technology
Publications Group
P.O. Box 73
Lexington, MA 02173-0073

ABSTRACT

An important component of any automated image analysis system is the detection and classification of objects. In this report, we consider the first of these problems where the specific goal is to detect anomalous areas (e.g., man-made objects) in textured backgrounds such as trees, grass, and fields of aerial photographs. ^{Their} Our detection algorithm relies on a significance test which adapts itself to the changing background in such a way that a constant false alarm rate is maintained. Furthermore, this test has a potentially practical implementation since it can be expressed in terms of the residuals of an adaptive two-dimensional linear predictor. The algorithm is demonstrated with both synthetic and real-world images.

Distribution/	
Availability Codes	
List	Avail and/or Special
A	



TABLE OF CONTENTS

ABSTRACT	111
1. INTRODUCTION	1
2. SIGNIFICANCE TESTING	3
3. DETECTION BASED ON LINEAR PREDICTION RESIDUALS	8
3.1 The Relationship of Significance Testing with Linear Prediction	8
3.2 Approximations Based on Image Modeling	11
3.3 The Question of Directionality	15
4. CONSTANT FALSE ALARM RATE DETECTION	17
5. ADAPTIVE ESTIMATION	21
6. IMPLEMENTATION OF THE DETECTION ALGORITHM	27
7. EXAMPLES	31
8. COMMENTS	57
ACKNOWLEDGMENTS	76
REFERENCES	77



1. INTRODUCTION

The problem of detecting small regions of an image which differ from their surroundings is of considerable interest in areas such as optical aerial reconnaissance, radar and infrared image analysis, and medical diagnosis through imagery. In the application of aerial reconnaissance, detection of such "anomalous" areas (or objects) of an image is often the first step in image analysis systems which perform automatic classification of man-made objects [1]. In this report, we shall address the particular problem of detecting objects in natural terrain (i.e., textured backgrounds) such as trees, grass, and fields of aerial photographs. We shall view an object as an area of an image with different second-order statistical properties from the surrounding area or background. Furthermore, we assume that the object's statistics are generally unknown (it is desired to detect broad classes of objects), but that the background statistics may be known or can be estimated.

Usually, in detection theory, the object (or signal) is added to the background (or noise), and filtering procedures are well-established for increasing the signal-to-noise ratio [2]. In image processing, however, the object pixels replace the background pixels. Motivated by this observation and the assumption that the object's statistics are unknown, we decide on the presence or absence of an object through significance testing [3]. In applying significance testing, we shall assume the background is characterized by a Gaussian probability density function. If a set of pixels falls in a critical region of this density function (i.e., low regions of probability), we reject the hypothesis that the pixels form part of the background.

To avoid estimating and inverting a large covariance matrix required in such a test, we impose structure on the background through modeling. As a first step toward this end, our significance test is expressed in terms of error residuals of two-dimensional (2-D) linear prediction. More specifically, the test, under a Gaussian assumption, first requires determining the error residuals from optimally predicting (in a least squares sense) each pixel from a linear combination of its neighboring pixels. The error residuals are then summed over a small area, suitably normalized, and finally compared to a threshold. Since a 2-D prediction filter is associated with our significance test, we can interpret this as representing the background by a 2-D autoregressive model [4]. The parameters of this autoregressive model are therefore assumed known or estimatable from the background. Furthermore, when the order of the model is fixed and small, an approximate, but a computationally practical implementation of the test results. It is interesting to note that the linear prediction residual has been used in a number of detection problems such as seismic event detection [5,6] and in detecting pitch in speech waveforms [7]. In the context of image processing, the prediction residual has been used successfully in, for example, segmenting textured images [4].

Since the background characteristics of an image are not stationary, i.e., are changing with position, in order to guarantee a constant false alarm rate (CFAR) [8] over the entire image, we must vary the threshold within our significance test as a function of the pixel position. We shall show that through our prediction interpretation of the test, a simple adaptive thresholding procedure arises which yields CFAR detection. The nonstationarity of an image also implies that we must adaptively esti-

mate the autoregressive model parameters at each pixel. The particular procedure employed is based on the 2-D covariance method of linear prediction [9] which is amenable to a recursive computation. In summary then, through adaptive estimation and thresholding, our significance test adapts itself to the changing background statistics to guarantee CFAR detection.

In the final pages of this report, the algorithm is successfully demonstrated through automatically detecting small-extent objects in real and synthetic images with varying textured backgrounds. Our real images are extracted from aerial photographs obtained from the Rome Air Force Development Center (RADC) data base. In these examples, we explore different approximations to the exact significance test. In particular, first and second quadrant prediction filters, averages of such filters, and noncausal prediction filters are investigated.

2. SIGNIFICANCE TESTING

The problem of object detection in images is viewed as one of finding small areas in an image whose statistical properties do not match those of the surrounding area or background. Essentially, we wish to determine whether a set of pixels under examination represent purely background or whether they contain partly or all object. The area of statistics that addresses such questions is called significance testing [3]. The basic idea is illustrated in Figure 1. A measurement is made of some random phenomenon characterized by probability density $p(\underline{x})$. A critical region C of low probability (α) is chosen corresponding to unlikely events. If measurements fall in the critical region, we reject the hypothesis that the measurements really belong to the density $p(\underline{x})$. The significance




Fig. 1. Probability density with critical region C.

level of the test is determined by the probability of events (α) in the critical region. For example, if this probability $\alpha=.05$, then the significance test is said to be at the 5% level.

In our case, the probability density is that of the background. If a set of measurements falls in the critical region, then we reject the hypothesis that the pixels under consideration form a part of the background; that is, we decide that they represent (at least partly) an object. More specifically, given an image and small region S at (n,m) , let the image points in S be denoted by (see Figure 2)

$$\underline{x} = \{x_1, x_2, \dots, x_N\}^T \quad (1)$$

We want to decide whether the points in S correspond to a background random field with probability density $p(\underline{x})$ (i.e., S contains just background) or whether S contains something other than the background random field (object possibly present). We want to do this for all (n,m) .

Thus, we must determine from the background probability density function (which we assume is known or estimatable) a critical region C of small probability (α) which is the level of significance. A critical region C can be defined by

$$p(\underline{x}) < \lambda \quad (2a)$$

so that the relation between λ and the level of significance α is given by

$$\int_R p(\underline{x}) d\underline{x} = \alpha \quad (2b)$$

where the region R contains only the points \underline{x} which satisfy $p(\underline{x}) < \lambda$.

Thus, in our significance test, if (2a) is satisfied, we decide "more than just background"; otherwise, we decide "background only". We repeat this for every (n,m) .

Note that the level of significance equals the probability of false




Fig. 2. An image and small region S .

alarm [2], i.e., the probability that we say an object is present given that we have pure background. Ideally, we wish to make the false-alarm probability as small as possible while making the probability of detection as large as possible (i.e., the probability of saying a target is present when a target is indeed present). However, since we assume no known statistics about the target, we cannot in this problem define such a probability. Consequently, we are forced to compute the probability of detection empirically.

Let us now suppose that the background corresponds to a Gaussian random process with mean $\underline{m} = E[\underline{x}]$ and covariance $K = E[(\underline{x} - \underline{m})(\underline{x} - \underline{m})^T]$. Then

$$p(\underline{x}) = \frac{1}{(2\pi)^{N/2} |K|^{1/2}} \exp \left[-\frac{1}{2} (\underline{x} - \underline{m})^T K^{-1} (\underline{x} - \underline{m}) \right] \quad (3)$$

We shall assume that the background is generally nonstationary, so that the covariance matrix K has no special structure. Nevertheless, we shall assume for the present that this matrix is known or can be estimated. Then, taking the logarithm of $p(\underline{x}) < \lambda$, from (3), our significance test becomes:

$$(\underline{x} - \underline{m})^T K^{-1} (\underline{x} - \underline{m}) > f(K, \lambda) \quad (4a)$$

with

$$f(K, \lambda) = \ln[(2\pi)^N |K|] - 2 \ln \lambda \quad (4b)$$

where the function $f(K, \lambda)$ is considered a threshold, corresponding to a certain probability of false alarm.

In the next section, we show that this test can be expressed in terms of the residuals of an adaptive 2-D linear predictor. The reasons for

this alternative formulation of the test are multi-fold. Perhaps the most important is to avoid the requirement of estimating and inverting a large covariance matrix. For example, if S is of extent 4×4 , then K is of extent 16×16 . We shall see that the residual error interpretation, along with imposing an autoregressive model of the background, leads to an approximate implementation of the test, requiring far fewer correlation coefficients. Furthermore, the true test involves a "one-shot" approach to the problem; i.e., a covariance matrix is estimated and used in the thresholding operation. The prediction approach, on the other hand, takes apart the true test into a number of components. Such a decomposition allows for both an alternative intuitive perspective of the algorithm and also a means of "twiddling" the various components to improve the test. We shall also see in section 4 that this decomposition leads to a choice of λ for guaranteeing CFAR detection with a nonstationary background.

3. DETECTION BASED ON LINEAR PREDICTION RESIDUALS

We now wish to show that the significance test of the previous section can be expressed in terms of the error residuals in optimally predicting each sample of our small region S by certain linear combinations of samples within S . This interpretation leads to a number of useful approximations to the true test when the image background is modeled by an autoregressive process.

3.1 The Relationship of Significance Testing with Linear Prediction

Our connection relies on the fact that since the background covariance matrix K is symmetric and positive definite it can be uniquely fac-

tored in terms of upper and lower triangular and diagonals' matrices [10].

In particular, we have

$$K = LDL^T \quad (5)$$

where L is lower triangular with one's along its diagonal, where D is a diagonal matrix, and where T denotes transpose. Substituting (5) into (4), we have

$$\begin{aligned} (\underline{x-m})^T K^{-1} (\underline{x-m}) &= (\underline{x-m})^T (LDL^T)^{-1} (\underline{x-m}) \\ &= (\underline{x-m})^T (L^{-1})^T D^{-1} L^{-1} (\underline{x-m}) \\ &= \underline{e}^T D^{-1} \underline{e} > f(K, \lambda) \end{aligned} \quad (6a)$$

where

$$\underline{e} = L^{-1} (\underline{x-m}) \quad (6b)$$

It is straightforward to show that since L is lower triangular with unit diagonal, L^{-1} has the same property and thus (6b) represents a causal transformation of the vector $\underline{x-m}$ [10]. That is, each e_k is a function of x_{l-m_l} for $l \leq k$. Furthermore, it can be shown that this transformation corresponds to successive orders of linear prediction where the diagonal elements of D are the prediction error variances [10]. That is, each row of L^{-1} represents the coefficients required in optimally predicting each element of the zero-mean vector $\underline{x-m}$ from a linear combination of its previous values. Since the covariance matrix K is that of the background, the coefficients of L^{-1} correspond to optimal prediction of the background and not object areas. This prediction concept is illustrated in Figure 3 for a zero-mean process. More specifically, an element




Fig. 3. Prediction of x_k from its previous values.

e_k of \underline{e} in (6b) is given by

$$e_k = (-\underline{a}_k^T, 1) \left(\begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_k \end{bmatrix} - \begin{bmatrix} m_1 \\ m_2 \\ \vdots \\ m_k \end{bmatrix} \right) \quad (7)$$

where \underline{a}_k is the vector of coefficients for optimally predicting $x_k - m_k$ from the previous values of $\underline{x} - \underline{m}$. The diagonal elements of D are given by $\sigma_k^2 = \text{var}(e_k)$. Note that the e_k 's are uncorrelated, i.e., they form a white process when the pixels being predicted (and doing the prediction) are background pixels [10].

Returning now to our significance test, we have from (6a),

$$\underline{e}^T D^{-1} \underline{e} = \sum_{k=1}^N \frac{e_k^2}{\sigma_k^2} > f(K, \lambda) \quad (8)$$

where σ_k^2 is the prediction error variance associated with predicting a background value x_k with its mean subtracted. Thus, the significance test involves first forming the prediction residuals e_k over S (from growing predictors) and normalizing e_k^2 with the corresponding prediction error variance σ_k^2 . These normalized residuals are summed over S and then compared to the threshold $f(K, \lambda)$.

3.2 Approximations Based on Image Modeling

We saw in the previous section that samples used in prediction are "causally" related to each x_k being predicted. Furthermore, we would

like to say that each x_k is predicted from its "past" values which fall within S. The notion of causality is thus based on a definition of past. For any point (n_0, m_0) , we define the past to be the set of points

$$\{(n, m) \mid n = n_0, m < m_0; n < n_0, -\infty \leq m \leq \infty\} \quad (9)$$

which are illustrated in Figure 4. As a matter of notation, if (n_1, m_1) is in the past of (n_2, m_2) , we denote this by $(n_1, m_1) < (n_2, m_2)$.




Fig. 4. Definition of past.

Let us now suppose that each pixel of the background image is known to be linearly related to its past. Specifically, let's suppose that the background is generated by an autoregressive model of the form:

$$x(n,m) = \sum_j \sum_k a(n,m;j,k) x(n-j,m-k) + \sigma(n,m)w(n,m) \quad (j,k) > (0,0) \quad (10)$$

where $w(n,m)$ is white Gaussian noise with unit variance and where the model coefficients and the variance of the driving function may vary in space (i.e., we assume the background is nonstationary). We shall refer to such a model as a nonstationary causal autoregressive model.

Let us further suppose that the background follows a causal autoregressive model of finite order. For example, suppose the background follows a third order model of the form:

$$x(n,m) = a(n,m;1,0) x(n-1,m) + a(n,m;0,1) x(n,m-1) + a(n,m;1,1) x(n-1,m-1) + \sigma(n,m)w(n,m) \quad (11)$$

Since the $a(n,m;j,k)$'s and the impulse response associated with this model have a 1st-quadrant region of support, we refer to (11) as a 1st-quadrant model.

Consider now using in the significance test of (8) a fixed-order predictor with the same region of support as the $a(n,m;j,k)$'s in (11), rather than a growing predictor. Then, except for the L-shaped boundary elements, illustrated in Figure 5, each element of e_k in (8) equals the prediction error from the growing predictor. That is, outside of the boundary, the remaining coefficients under the (growing) prediction mask of Figure 4 are zero, and thus do not effect the prediction. Note that with the fixed-order predictor, elements outside of our region S will now be used to predict the boundary elements.




Fig. 5. Representation of a first-quadrant space-varying fixed-order predictor.

More generally, the background random field can be well represented by a causal autoregressive model of arbitrarily large, but finite order [11]. Consequently, there exists generally a larger boundary region (i.e., larger than that of Figure 5) where the fixed-order prediction error deviates from that of the prediction error corresponding to the growing predictor of the true significance test. It is important to observe that with a fixed-order predictor, the number of required correlation coefficients now depends on the order of the fixed predictor and not on the size of S , as in the true significance test. We shall return to this issue in section 4.

Finally, an approximate significance test can be written as:

$$\sum_{k=1}^N \tilde{e}_k^2 / \tilde{\sigma}_k^2 > f(K, \lambda) \quad (12)$$

where \tilde{e}_k and $\tilde{\sigma}_k^2$ are the prediction error and prediction error variance, respectively, associated with a fixed-order prediction of each element of S . Note that although the predictor is fixed in order, it does vary in space over S (as illustrated in Figure 5), since we have assumed the background to be generally nonstationary.

3.3 The Question of Directionality

It is curious that although the significance test was derived with no imposed directionality, the linear predictor which results is causal. This apparent contradiction can be resolved by noting that the causality of the growing predictor in (6) arises only because of our way of ordering the samples of S .

For example, suppose that we order the samples in reverse order as illustrated in Figure 6. With this ordering, the predictors become "anti-causal". As before, we can approximate these growing predictors by fixed-order predictors.

More generally, we can assume any ordering of the samples in S (e.g., diagonal, random, etc.) Our choice in an approximate significance test (using a fixed-order predictor) ultimately depends on how close the background process matches the assumed pixel relationship. Since textures in images appear to have no directionality, some sort of noncausal prediction mask may be most appropriate. Alternately, we may attempt to remove directionality imposed by a fixed-order causal predictor by averaging many predictors of different directionalities as, for example, 1st, 2nd, 3rd, and 4th-quadrant predictors.




Fig. 6. "Anti-causal" prediction of x_k .

4. CONSTANT FALSE ALARM RATE DETECTION

We now consider determining λ in our threshold $f(K, \lambda)$ for a constant false alarm rate. That is, since the background covariance matrix K will change as our region S sweeps over the image, we must twiddle λ in accordance with the changes in K so that the integral in (2) remains constant. In this section, we derive this required functional form of λ . Our result relies on the orthogonalization of the elements of \underline{x} through the matrix decomposition (5).

Without loss of generality, we shall assume a zero-mean process. Then from (2), (3), and (6), we have:

$$\begin{aligned} \alpha &= \int_{p(\underline{x}) < \lambda} \frac{1}{(2\pi)^{N/2} |K|^{1/2}} \exp \left[-\frac{1}{2} \underline{x}^T K^{-1} \underline{x} \right] d\underline{x} \\ &= \int_{p(\underline{x}) < \lambda} \frac{1}{(2\pi)^{N/2} |K|^{1/2}} \exp \left[-\frac{1}{2} (L^{-1} \underline{x})^T D^{-1} (L^{-1} \underline{x}) \right] d\underline{x} \end{aligned} \quad (13)$$

Now, let

$$\underline{e} = L^{-1} \underline{x} \quad (14)$$

so that, since L^{-1} has unit diagonals, using the method of Jacobians [12], we have,

$$d\underline{e} = d\underline{x} \quad (15)$$

Thus, substituting (14) and (15) into (13), we obtain

$$\int_{p(L\underline{e}) < \lambda} \frac{1}{(2\pi)^{N/2} |K|^{1/2}} \exp \left[-\frac{1}{2} \underline{e}^T D^{-1} \underline{e} \right] d\underline{e} = \alpha \quad (16)$$

Furthermore, we have (since $D^{-1/2}$ is diagonal):

$$\begin{aligned}
 & \int_{p(\underline{L}\underline{e}) < \lambda} \frac{1}{(2\pi)^{N/2} |\underline{K}|^{1/2}} \exp \left[-\frac{1}{2} (\underline{D}^{-1/2} \underline{e})^T (\underline{D}^{-1/2} \underline{e}) \right] d\underline{e} \\
 &= \int_{p(\underline{L}\underline{D}^{1/2} \hat{\underline{e}}) < \lambda} \frac{1}{(2\pi)^{N/2} |\underline{K}|^{1/2}} \exp \left[-\frac{1}{2} \hat{\underline{e}}^T \hat{\underline{e}} \right] |\underline{D}|^{1/2} d\hat{\underline{e}} \\
 & \qquad \qquad \qquad = \alpha
 \end{aligned} \tag{17}$$

where we have used the substitutions:

$$\hat{\underline{e}} = \underline{D}^{-1/2} \underline{e} \tag{18a}$$

$$d\hat{\underline{e}} = |\underline{D}|^{-1/2} d\underline{e} \tag{18b}$$

Noting that $|\underline{K}| = |\underline{D}|$ [3], we have from (18b),

$$\int_{p(\underline{L}\underline{D}^{1/2} \hat{\underline{e}}) < \lambda} \frac{1}{(2\pi)^{N/2}} \exp \left[-\frac{1}{2} \hat{\underline{e}}^T \hat{\underline{e}} \right] d\hat{\underline{e}} = \alpha \tag{19}$$

Observe that the integrand in (19) involves the orthogonal (i.e., white) elements e_k and does not depend on the statistics of \underline{x} . Furthermore, considering the limits of integration in (19) we see that the boundary of our (transformed) critical region is given by the equation:

$$p(\underline{L}\underline{D}^{1/2} \hat{\underline{e}}_0) = \lambda$$

which, from (2) and (6), can be expressed as

$$\begin{aligned} & \frac{1}{(2\pi)^{N/2} |K|^{1/2}} \exp \left[-\frac{1}{2} (LD^{1/2} \hat{\underline{e}}_0)^T L^{-1} D^{-1} L^{-1} (LD^{1/2} \hat{\underline{e}}_0) \right] \\ &= \frac{1}{(2\pi)^{N/2} |K|^{1/2}} \exp \left[-\frac{1}{2} \hat{\underline{e}}_0^T \hat{\underline{e}}_0 \right] = \lambda \end{aligned} \quad (20a)$$

or,

$$-\frac{1}{2} \hat{\underline{e}}_0^T \hat{\underline{e}}_0 = \ln[(2\pi)^{N/2} |K|^{1/2} \lambda] \quad (20b)$$

The transformation of \underline{x} to $\hat{\underline{e}}$ through $D^{-1/2} L^{-1}$ and the corresponding boundary of integration are shown in Figure 7.

We see then from (20) and Figure 7 that to maintain a constant false alarm rate, we must have:

$$\ln[(2\pi)^{N/2} |K|^{1/2} \lambda] = \text{constant} \quad (21a)$$

or

$$\lambda = \frac{C}{(2\pi)^{N/2} |K|^{1/2}} \quad (21b)$$

Substituting (21) into (6), we obtain:

$$\begin{aligned} \underline{\underline{e}}^T D^{-1} \underline{\underline{e}} &> f(K, \lambda) \\ &= -\ln [(2\pi)^N |K|] - 2\ln \lambda \\ &= -\ln [(2\pi)^N |K|] + \ln [(2\pi)^N |K|] + \ln C^{-2} \\ &= \text{constant} \end{aligned} \quad (22)$$




Fig. 7. Pictorial representation of orthogonal transformation.

Finally, then, our approximate significance test in (12) based on a fixed-order predictor is given by

$$\sum_{k=1}^N \tilde{e}_k^2 / \tilde{\sigma}_k^2 > \text{constant} \quad (23)$$

Equation (23) lends itself to a simple intuitively pleasing interpretation. Let's suppose that the background follows our assumed fixed-order autoregressive model. Then generating \tilde{e}_k represents an attempt to first whiten the data. Normalization by the variance of the white residual gives equal weight to the generally nonstationary image pixels. Finally, we perform a significance test on sets of N (equally "important") white Gaussian samples of unit variance.

5. ADAPTIVE ESTIMATION

The significance test of the previous sections depends on knowledge of the background statistics; i.e., we need to know or estimate the coefficients of the assumed space-varying autoregressive model. However, in attempting this estimation, we encounter the uncertainty principle. That is, to obtain a reliable estimate of $a(n,m;j,k)$, we require stationarity over a "sufficiently large" window size. On the other hand, we assume statistics are generally changing everywhere in space.

To side-step this problem, we assume that the data is in fact stationary over the extent of what we shall refer to as the estimation window, $w_e(n,m)$. The location of the 2-D estimation window which slides over our data will be designated by the center index (n_0, m_0) as illus

trated in Figure 8. The model parameters associated with this window are defined to be those at (n_0, m_0) : $a(n_0, m_0; j, k)$.

The particular least squares estimation procedure we shall use is the covariance method of linear prediction [9]. The prediction error associated with the estimation window at spatial coordinates (n_0, m_0) is the error $e(n_0, m_0)$ in predicting the value $x(n_0, m_0)$ from its neighbors weighted by $a(n_0, m_0; j, k)$. Finally, the prediction error variance $\sigma^2(n_0, m_0)$ is given by the average squared prediction error under the estimation window at (n_0, m_0) , based on the coefficients $a(n_0, m_0; j, k)$.

For each pixel location (n_0, m_0) , we wish to estimate the set of model parameters $a(n_0, m_0; k, l)$ which vary in space. To do this, we assume that

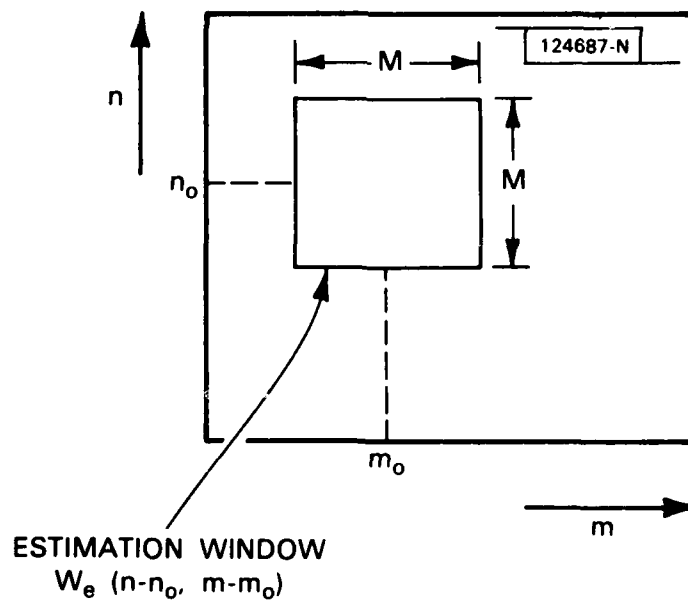


Fig. 8. Representation of the estimation window.

$x(n,m)$ is a 2-D random field stationary over each $w_e(n-n_0, m-m_0)$, and follows the model (10), but where the parameters $a(n,m;j,k)$ and $\sigma(n,m)$ are assumed constant with respect to n and m over the estimation window. Therefore, throughout the remainder of this section, we drop the spacial dependence and work with the model given by

$$x(n,m) = \sum_{(j,k) > (0,0)} a(j,k) x(n-j, m-k) + \sigma w(n,m)$$

We shall assume that the prediction coefficients $a(j,k)$ fall within a $(P \times Q)$ first-quadrant plane mask. For simplicity, we limit our derivations to this class of prediction masks, although it is clearly applicable to more general mask shapes, such as other quadrant masks and non-causal masks. Our objective is to estimate from $x(n,m)$ the model parameters $a(j,k)$ for $j=0,1,\dots,P-1$ and $k=0,1,\dots,Q-1$, with $j=k \neq 0$. Further, let's suppose that we have available $x(n,m)$ for $(n,m) \in [-P+n_1, n_2] \times [-Q+m_1, m_2]$ (see Figure 9). We then define the error $e(n,m)$ over the region I , given by $I = [n_1, n_2] \times [m_1, m_2]$, as

$$e(n,m) = x(n,m) - \sum_{\substack{j=0 \\ (j,k) \neq (0,0)}}^{P-1} \sum_{k=0}^{Q-1} a(j,k) x(n-j, m-k) \quad (n,m) \in I \quad (25)$$

Our goal becomes to minimize the sum of the squared errors given by

$$E[n_0, m_0] = \sum_{n=n_1}^{n_2} \sum_{m=m_1}^{m_2} e^2(n,m) \quad (26)$$

Note that the region I over which $e(n,m)$ is given is equivalent to the region under $w_e(n-n_0, m-m_0)$, but that in determining $e(n,m)$, we have used some data along the border of $w_e(n-n_0, m-m_0)$.




Fig. 9. Known data blocks used in 2-D least squares.

The approach we take is to transform the 2-D problem to a 1-D problem so that a 1-D least squares solution is applicable. Note, however, that we will still have solved the 2-D least squares problem. In particular, we wish to transform (26) into a 1-D error expression. To accomplish this transformation, we define the vectors $a[n_0, m_0]$ and σ by

$$\begin{array}{c}
 \begin{array}{c}
 a(n_0, m_0) = \begin{bmatrix}
 a(0,1) \\
 a(0,2) \\
 \vdots \\
 a(0,Q) \\
 \hline
 a(1,0) \\
 a(1,1) \\
 \vdots \\
 a(1,Q) \\
 \hline
 \vdots \\
 \hline
 a(P,0) \\
 a(P,1) \\
 \vdots \\
 a(P,Q)
 \end{bmatrix}
 \end{array}
 \end{array}
 \begin{array}{c}
 \uparrow \\
 (PQ-1) \\
 \downarrow
 \end{array}
 \begin{array}{c}
 \sigma = \begin{bmatrix}
 x(n_1, m_1) \\
 x(n_1, m_1+1) \\
 \vdots \\
 x(n_1, m_2) \\
 \hline
 x(n_1+1, m_1) \\
 x(n_1+1, m_1+1) \\
 \vdots \\
 x(n_1+1, m_2) \\
 \hline
 \vdots \\
 \hline
 x(n_2, m_1) \\
 x(n_2, m_1+1) \\
 \vdots \\
 x(n_2, m_2)
 \end{bmatrix}
 \end{array}
 \begin{array}{c}
 \uparrow \\
 (PQ-1) \\
 \downarrow
 \end{array}
 \quad (27)$$

and the matrix S by

$$\begin{array}{c}
 S = \begin{bmatrix}
 \overleftarrow{(PQ-1)} \quad \overrightarrow{} \\
 \hline
 A_0 \\
 \hline
 A_1 \\
 \hline
 \vdots \\
 \hline
 A_{M-1} \\
 \hline
 \end{bmatrix}
 \begin{array}{c}
 \uparrow \\
 M^2 \\
 \downarrow
 \end{array}
 \end{array}
 \quad (28a)$$

where,

$$A_j = \begin{matrix} \xleftarrow{\quad PQ-1 \quad} \xrightarrow{\quad} \\ \left[\begin{array}{cc} [x(n_1+j-0, m_1-1) \dots x(n_1+j-0, m_1+Q)] & \dots [x(n_1+j-P, m_1-0) \dots x(n_1+j-P, m_1-Q)] \\ [x(n_1+j-0, m_1+1-1) \dots x(n_1+j-0, m_1+1+Q)] & \dots [x(n_1+j-P, m_1+1-0) \dots x(n_1+j-P, m_1+1-Q)] \\ [x(n_1+j-0, m_2-1) \dots x(n_1+j-0, m_2-Q)] & \dots [x(n_1+j-P, m_2-0) \dots x(n_1+j-P, m_2-Q)] \end{array} \right] \end{matrix} \quad (28b)$$

and where we have assumed the data segment I to be of extent $M \times M$. Note that σ is a vector consisting of the concatenation of the rows of $x(n, m)$ over I, $a[n_0, m_0]$ is a vector consisting of the concatenation of the rows of $a(j, k)$ for $(j, k) \in [0, P] \times [0, Q]$ with $(j, k) \neq (0, 0)$, and S is a matrix which consists of the concatenation of rows of various subsequences of the known $x(n, m)$ required in predicting each value of $x(n, m)$ over I.

Therefore, we can write (26) as:

$$\begin{aligned} E[n_0, m_0] &= \sum_{n=n_1}^{n_2} \sum_{m=m_1}^{m_2} e^2(n, m) \\ &= (Sa[n_0, m_0] - \sigma)^T (Sa[n_0, m_0] - \sigma) \end{aligned} \quad (29)$$

We then write the solution to minimizing (29) with respect to $a[n_0, m_0]$ as [9]:

$$a[n_0, m_0] = R^{-1} S^T \sigma \quad (30a)$$

where

$$R = S^T S \quad (30b)$$

Note that the matrix R is of extent $(PQ-1) \times (PQ-1)$ so that computation required in its inversion is dependent on the model order.

In particular, since R is generally not Toeplitz, its inversion will require on the order of $(PQ-1)^3$ operations. Thus assuming $P, Q \ll M$, the bulk of the computation is embedded within forming $R = S^T S$ which requires on the order of M^4 operations.

This estimation is then carried out at each pixel. An alternative to this direct estimation is to accomplish the estimation recursively [13]. However, this appears to be a viable alternative only when the estimation window size is less than the model order [14]; i.e., the matrix required to be inverted at each pixel is on the order of $M \times M$. We are currently investigating methods to reduce this computation.

In either case, we obtain a parameter set at each pixel which represents an estimate of the model parameters of the changing background, required in our prediction procedure. Finally, it is straightforward to show from (29) and (30) that the estimate of the prediction error variance given by the average squared prediction error under each estimation window can be expressed by

$$\hat{\sigma}^2(n_0, m_0) = (\sigma^T \sigma - a[n_0, m_0]^T (S^T S) a[n_0, m_0]) N^{-2} \quad (31)$$

5. IMPLEMENTATION OF THE DETECTION ALGORITHM

We are now ready to merge the results of the previous sections to form an implementation of our target detection algorithm. From our coefficient estimates (30), we compute the prediction error function $\hat{e}^2(n, m)$

based on a fixed-order, but space-varying prediction model. Then with $\tilde{\sigma}^2(n,m)$ in (31), we can write a 2-D version of the approximate significance test of (23), as:

$$\sum_{k,l \in S} \tilde{e}^2(k,l) / \tilde{\sigma}^2(k,l) > \text{constant} \quad (32)$$

where we can think of the indices k and l as running over different regions S which sweep over the images.

Equivalently, we can consider generating the statistic in (32) at each spatial location (n,m) of an image by convolving an $L \times L = N$ - point smoothing window, $w_s(n,m)$ (which falls over our region S) with the normalized prediction error to create a new smoothed function $E_s(n,m)$:

$$E_s(n,m) = q(n,m) ** w_s(n,m) \quad (33a)$$

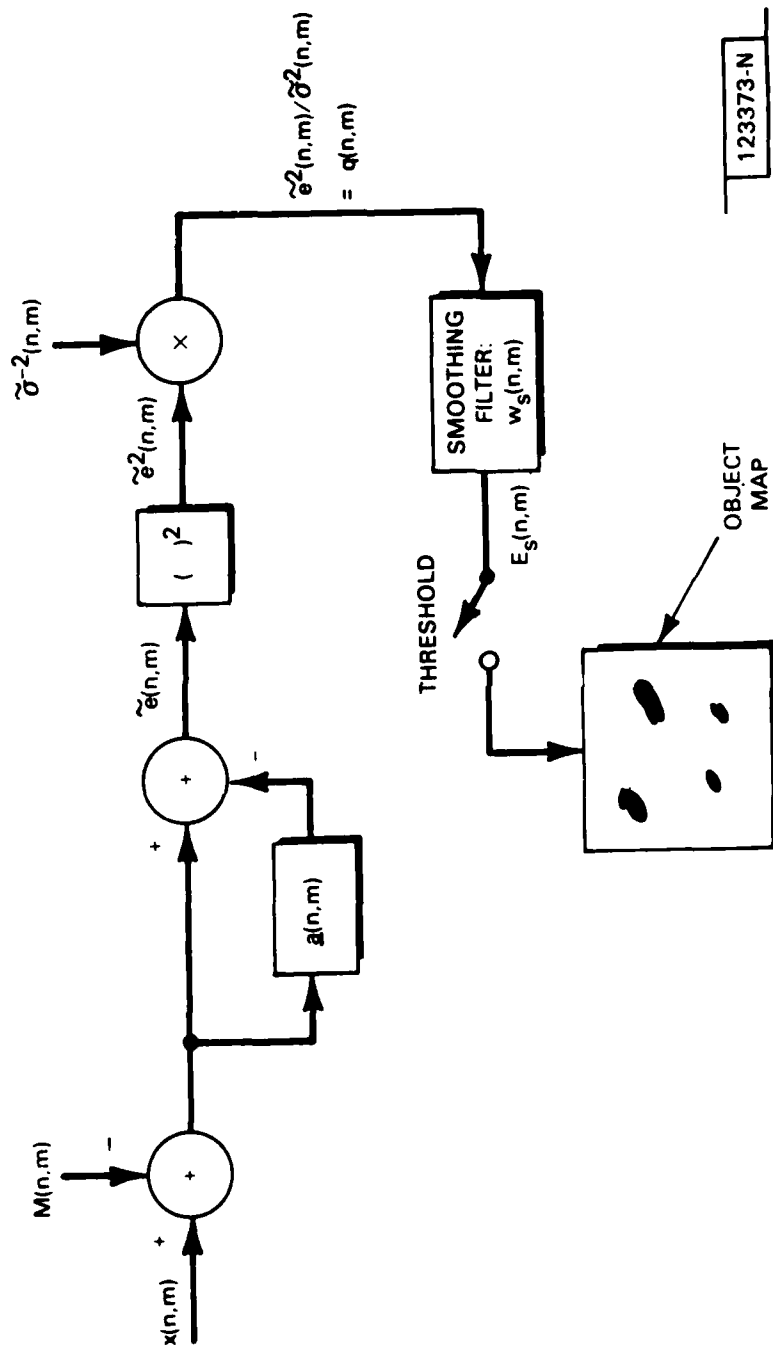
where

$$q(n,m) = \tilde{e}^2(n,m) / \tilde{\sigma}^2(n,m) \quad (33b)$$

In the estimation of the model parameters, the estimation window $w_e(n,m)$ should be small enough to preserve approximate stationarity, but large enough to obtain a reliable estimate of the required correlation coefficients. The estimation window must also be large enough so that anomalies (i.e., objects) do not badly corrupt the correlation estimates.

The smoothing window should be small enough so that small-extent objects are not overwhelmed by background in the significance test. However, it may also need be large enough so boundary effects in our finite-order model assumption do not play a significant role.

The overall detection algorithm based on the approximate significance test is illustrated in Figure 10. The first operation subtracts an estimate of the local mean of $x(n,m)$ which is computed by averaging $x(n,m)$ under $w_e(n,m)$. (Recall that our significance test requires a zero-mean random field.) Under the estimation window, a local covariance matrix R as defined in (30b) is computed. R is then used to find the coefficient vector $a[n,m]$ and the prediction error variance $\hat{\sigma}^2(n,m)$ in (30) and (31), respectively, which are required to compute the normalized prediction error, $q(n,m)$. Finally, $q(n,m)$ is convolved with the smoothing window $w_s(n,m)$ and compared to a threshold.



123373-N

Fig. 10. The detection algorithm.

7. EXAMPLES

In this section, we present a number of examples based on the detection algorithm developed in the previous sections. Throughout this section, we have chosen the estimation window $w_e(n,m)$ to be of size 10×10 pixels which we assume is "sufficiently" larger than the size of most objects. This assumption can be justified through our empirical observation that in most cases the coefficient change function (CCF) [14] of our processed images is relatively flat; i.e., the objects' presence appears to not adversely affect estimation of background statistics. We also assume a 10×10 window is large enough to obtain a good estimate of the correlation coefficients required in estimating $a[n,m]$ and $\sigma^2(n,m)$, but also small enough to maintain approximate stationarity. Of course, however, this assumption breaks down at region boundaries.

In our first examples, we consider computer generated 1-D and 2-D signals determined by exciting all-pole filters with white noise. We then analyze progressively more complicated real images which we have obtained from the Rome Air Force Development Center (RADC) data base.


EXAMPLE 1:

Consider a sequence $x(n)$ of the form:


$$x(n) = 0.95 x(n - 1) + w(n) \quad (34)$$

where $w(n)$ is zero-mean white noise. A sample function of $x(n)$ is shown in Figure 11.a and a 1-point "object" at $n = 64$ is shown in Figure 11.b. The (single) coefficient estimate was based on a 16-point estimation window. Figure 11.c shows the squared prediction error \hat{e}_k^2 . The object is clearly detected.


Consider a second sequence depicted in Figure 12.a of the form in (34) created with a different white-noise input. A four-point object has been implanted at locations $n = 90, 91, 92$ and 93 . As before the (single) coefficient estimate was based on a 16-point estimation window. The squared prediction error, illustrated in Figure 12.b, gives a clear indication of the object.



(a)




(b)




(c)

Fig. 11. Detection of 1-point object in Example 1: (a) 1-D random sequence; (b) random sequence with object; (c) squared prediction error.



(a)



(b)


Fig. 12. Detection of 4-point object in Example 1: (a) 1-D random sequence with object; (b) squared prediction error.

EXAMPLE 2:


Figure 13.a depicts a 1-D slice of an aerial photograph consisting of a grove of trees. In Figure 13.b, we have implanted a one-point object at $n = 64$. In this experiment, a six-parameter non-causal model was assumed:

$$x(n) = \sum_{\substack{k = -3 \\ k \neq 0}}^3 a(k)x(n-k) + w(n) \quad (35)$$


where $w(n)$ is white noise. The estimation window was chosen at ten points in duration. The squared prediction error, shown in Figure 13.c, clearly picks out the object. For reference, Figure 13.d depicts the squared prediction error without the one-point object. A second object and its corresponding squared prediction error are shown in Figures 14.a and 14.b, respectively. The estimation window is sixteen points in duration and a two-point noncausal model is assumed.




(a)



(b)




(c)




(d)

Fig. 13. Detection of 1-point object in Example 2: (a) slice of trees; (b) slice of trees with object; (c) squared prediction error with object; (d) squared prediction error without object.



(a)



(b)

Fig. 14. Detection of 4-point object in Example 2: (a) slice of trees with object; (b) squared prediction error.

EXAMPLE 3:

Consider a 2-D sequence generated by the particular 2-D difference equation of the form:

$$\begin{aligned} x(n,m) = & a(0,1)x(n-0,m-1)+a(1,0)x(n-1,m-0) \\ & +a(1,1)x(n-1,m-1)+w(n,m) \end{aligned} \quad (36)$$

The background sequence (64x64 pixels in extent) was generated with coefficients $a(0,1)=0.1$, $a(1,0)=-0.9$ and $a(1,1)=0.1$. Four objects were implanted within the image, all of a constant level, but with a variance about equal to that of the background. Moreover, the size and level of the anomalies were chosen to be visually difficult to detect from the background (see Figure 15.a). The model assumed in the estimation procedure is given by the generating process (36).

The 3-D perspective of the squared prediction error $\tilde{e}^2(n,m)$ is given in Figure 16. All four objects are clearly detected and even the two closely spaced objects are resolved. This same function, along with the smoothed $\tilde{e}^2(n,m)$ (a 3x3 smoothing window, $w_g(n,m)$, was applied in this example) are illustrated in Figures 15.b and 15.c after thresholding. In Figure 15.d is shown the prediction error variance and in Figures 15.e and 15.f the smoothed normalized prediction error -- both appropriately thresholded.

Note that two different thresholds are applied to the smoothed normalized prediction error. The first resolves three of the four objects -- the second resolves all four objects, but introduces false alarms. This is due to the inaccuracies of the estimate of the prediction error variance which is illustrated in Figure 15.d. Ideally, since the background

is stationary, the estimated prediction error variance should be flat. However, as seen in Figure 15.d, the estimate actually peaks in the region of objects — contrary to what we would hope to happen. We have encountered in this synthetic example, perhaps, what is a fundamental limitation in measuring the background prediction error variance: the presence of objects can (falsely) increase the background residual variance. With apriori knowledge that the background prediction error variance is constant, we were able to improve detection.

123378 R




Fig. 15. Detection of objects in test image for Example 3.
(a) test image with four objects, (b) prediction error,
(c) smoothed prediction error, (d) prediction error variance,
(e) smoothed normalized prediction error (high threshold),
(f) smoothed normalized prediction error (low threshold).

124963-S




Fig. 16. 3-D perspective of prediction error for Example 3.



EXAMPLE 4:

Figure 17.a depicts a 64x64-pixel RADC image in which two 2x2-pixel synthetic objects (of constant level) have been implanted. This image was created by a 64-to-1 downsampling and smoothing of the original image. The assumed background model is the same three-parameter model used in the previous example in (36). Figures 17.b-17.e depict the prediction error, prediction error variance, and smoothed normalized prediction error (a 4x4 smoother, $w_s(n,m)$, was applied), respectively. The processed part of the image is given within the boxed area. Note that normalization of the prediction error in this case (unlike the previous example) has helped in bringing out the object from the more busy field background.




Fig. 17. Detection of objects in RADC image for Example 4. (a) image with two synthetic objects, (b) prediction error, (c) prediction error variance, (d) smoothed normalized prediction error (high threshold), (e) smoothed normalized prediction error (low threshold).

EXAMPLE 5:

Figure 18.a depicts a 64x64-pixel RADC image in which two 3x3-pixel synthetic objects (of constant level) have been implanted. This field-tree image was created by a 64-to-1 downsampling and smoothing of the original image. In our first attempt to detect the two synthetic objects, the three-parameter model of (36) was assumed. Although the object in field background was easily detected, the object in tree background was not detected, even with normalization by the prediction error variance.

Consequently, in our second attempt at detection, we assumed a twelve-parameter non-symmetric half-plane [11] autoregressive model. This model is more general and thus more likely to accurately model the background [11]. Figures 18.b-18.e depict the prediction error, prediction error variance, and smoothed normalized prediction error (a 4×4 $w_s(n,m)$ was applied), respectively. Because of the computational intensity with a twelve-parameter model, only the designated region was processed. Note that normalization of the prediction error has helped significantly in bringing out from the background the object embedded within the trees.






Fig. 18. Detection of objects in RADC image for Example 5. (a) image with two synthetic objects, (b) prediction error, (c) prediction error variance, (d) smoothed normalized prediction error (high threshold), (e) smoothed normalized prediction error (low threshold).

EXAMPLE 6:

Consider the RADC image displayed in Figure 19.a. This image consists of 128x128 pixels and was created by a 16-to-1 downsampling and smoothing of the original image. The assumed background model is the same three-parameter model used in example 3 in (36). Figures 19.b-19.d also depicts the prediction error, smoothed prediction error and smoothed normalized prediction error — suitably thresholded. As in our synthetic example, the smoothed prediction error without normalization yields fewer detected objects (which may or may not be considered false alarms) than the smoothed normalized prediction error. This happens probably because the background variance appears reasonably constant throughout the image. The objects, however, can potentially introduce a false increase in the local variance, as illustrated in Figure 19.e which shows the thresholded prediction error variance.



123382 R



(a)



(b)



(c)



(d)



(e)

Fig. 19. Comparison of smoothed prediction error and smoothed normalized prediction error for Example 6. (a) RADC image, (b) smoothed normalized prediction error, (c) prediction error, (d) smoothed prediction error, (e) prediction error variance.

PREVIOUS PAGE
IS BLANK



EXAMPLE 7:

Consider the RADC image displayed in Figure 20.a. This image consists of 128x128 pixels and was created by a 16-to-1 downsampling and smoothing of the original image. As in the previous example, a three-parameter autoregressive model is assumed. Figure 20 makes the same comparisons among the various residuals as made in Example 6.






Fig. 20. Comparison of smoothed prediction error and smoothed normalized prediction error of Example 7. (a) RADC image, (b) smoothed normalized prediction error, (c) prediction error, (d) smoothed prediction error, (e) prediction error variance.

PREVIOUS PAGE
IS BLANK

EXAMPLE 8:


Consider the RADC image displayed in Figure 21.a. This image consists of 128x128 pixels and was created by an 64-to-1 downsampling and smoothing of the original image. In this example, we consider a first-quadrant causal, second-quadrant causal autoregressive model, and average of the two. This average represents an attempt to do away with directionality of the approximate significance test.

Figures 21 and 22 illustrate the results with first-quadrant (three-parameter) and second-quadrant (three-parameter) prediction masks, respectively. Figure 23 summarizes our results by depicting the smoothed normalized prediction errors and their average. Note that the individual smoothed normalized prediction errors do well in detecting most objects, while the average appears to deteriorate the performance.

Two additional experiments that were performed with this data are shown in Figures 24 and 25. In Figure 24, we show a thresholded version of the coefficient change function (CCF) [14] corresponding to the predictor of Figure 21. The CCF bears little resemblance to our prediction errors. Moreover, due to the large estimation window (i.e., 10x10 pixels), this function is small everywhere — reflecting little sample-to-sample change in the coefficient estimates. Finally, in Figure 25, we depict the smoothed noncausal normalized prediction error. The noncausal prediction mask is an eight-point nearest neighbor mask. The results on this image and others (e.g., example 6) are encouraging, but appear to do no better (and perhaps worse) than the causal prediction masks.



(a)



(b)



(c)



(d)

Fig. 21. Comparison of prediction error and smoothed normalized prediction error with first-quadrant mask for Example 8. (a) RADAR image, (b) prediction error variance, (c) prediction error, (d) smoothed normalized prediction error.




Fig. 22. Comparison of prediction error and smoothed normalized prediction error with second-quadrant mask for Example 8. (a) RADC image, (b) prediction error variance, (c) prediction error, (d) smoothed normalized prediction error.




Fig. 23. Average of smoothed normalized prediction errors for Example 8. (a) RADC image, (b) smoothed normalized prediction error (1st quad), (c) smoothed normalized prediction error (2nd quad), (d) average of (b) and (c).



Fig. 24. CCF corresponding to Figure 21.




Fig. 25. Smoothed normalized noncausal prediction error for Example 8.



EXAMPLE 9:

Consider the 64x64 pixel RADC image in Figure 26 generated by down-sampling the original image by 16-to-1 with smoothing. This image is particularly interesting due to the presence of a radio tower in the lower right-hand corner of the image. Note that the top of the radio tower has been clearly detected.


It is interesting to observe that in examples 7, 8, and 9 normalization of the prediction errors helped detection and reduced false alarms by reducing the background variance in busy regions such as the tree and brush areas.



123389 R




(a)



(b)



(c)



(d)

Fig. 26. Comparison of prediction error and smoothed normalized prediction error for Example 9. (a) RADC image, (b) prediction error variance. (c) prediction error, (d) smoothed normalized prediction error.

PREVIOUS PAGE
IS BLANK

8. COMMENTS

The examples of the previous section clearly illustrate the success of the 2-D prediction residual in object detection. However, there remain many unanswered questions. For example, we have not in this report implemented the true significance test. The difficulties involved in estimating a large covariance matrix makes this undesirable. Note also that the true test involves a "one-shot" approach to the problem. A covariance matrix is estimated and used in the thresholding operation. If the estimate is bad -- as it may be in object regions -- we have no chance to twiddle parameters interactively. The prediction approach, on the other hand, has taken the true test apart into a number of components, allowing individual twiddling of prediction parameters, normalization, etc. From this viewpoint, it may result that the true test is not a good standard and the approximations and their various modifications may yield better results.

Another unresolved area involves combining various approximations such as 1st and 2nd quadrant predictors to approach the true test. The few preliminary experiments with such combinations have not yielded consistent results. Nevertheless, there may exist some theoretical justification for combining different quadrant predictors. As we saw in section 3.3, by simply ordering samples in our region S in different ways, we obtain different implementations of the significance test. These implementations involve different directionalities, e.g., 1st or 2nd-quadrant growing prediction masks. Often the "bad" samples (i.e., samples not yielding prediction errors of the true test) of one implementation are the "good" samples of another. Consequently, it seems reasonable that there

exists a way to combine these different predictors to (theoretically) approach the true significance test.

Furthermore, since the significance test can be implemented with a (growing) predictor of any directionality, there does not exist inherent directionality within the test. The approximations, however, impose their own specific directionality. Recall, however, that the approximation can be made safely only when the background has associated with it certain directionality (e.g., 1st-quadrant or 2nd-quadrant models). Therefore, the validity of the directionality imposed by the approximation is directly linked to the assumption that the background has some sort of directionality associated with it -- perhaps, an unreasonable assumption.

In section 5, we only touched upon the estimation problem. Clearly, we might seek better estimates of the model parameters and of background variance which is used in normalization. An iterative technique is one possibility. On each iteration, we might estimate background statistics from pixels which do not include current object samples. Alternatively, we might fill in what we think are object areas with a signal predicted from the background. The former case raises questions about estimation with missing data.

An interesting characteristic of this detection algorithm is that lines and edges of regions tend to be suppressed, while anomalous areas are enhanced. It is also at lines and region boundaries where our modeling assumptions break down. A better understanding is needed of the response of the prediction error in such areas. For example, an "optimal" size of the region S may reduce the probability of false alarm in these regions and increase detections in anomalous areas.

Finally, we might consider introducing some general apriori knowledge about the objects of interest. A procedure for introducing such knowledge in a significance test remains to be understood. One possible approach might involve developing a "cross" between significance and hypothesis testing.

ACKNOWLEDGEMENTS

The author wishes to thank Dr. C.W. Therrien for supplying useful insights on significance testing and linear prediction. Dr. D. Dudgeon and Dr. T. Bially also engaged in many valuable discussions.

REFERENCES

1. W.K. Pratt, Digital Image Processing (Wiley, New York, 1978).
2. H.L. Van Trees, Detection, Estimation, and Modulation Theory (Wiley, New York, 1968).
3. A.W. Drake, Fundamentals of Applied Probability (McGraw-Hill, New York, 1967).
4. C.W. Therrien, "An Estimation-Theoretic Approach to Terrain Image Segmentation," Computer Graphics and Image Processing, to be published (February 1983).
5. J.F. Claerbout, "Detection of P-Waves from Weak Sources at Great Distances," Geophysics, 24, 197 (1964).
6. W.P. Dove and A.V. Oppenheim, "Event Location Using Recursive Least Squares Signal Processing", IEEE International Conference on Acoustics, Speech, and Signal Processing, Denver, Colorado, (April 1980), pp. 848-850.
7. D.T.L. Lee and M. Morf, "A Novel Innovations Approach Based Time-Domain Pitch Detector", IEEE International Conference on Acoustics, Speech, and Signal Processing, Denver, Colorado, (April 1980), pp. 40-43.
8. D.C. Schleher, Automatic Detection and Radar Data Processing (Artech House, Dedham, Massachusetts, 1980).
9. J.H. McClellan, "Parametric Signal Modeling Notes," Course Notes for Advanced Topics in Digital Signal Processing, (1981).
10. C.W. Therrien, "On the Relation Between Triangular Matrix Decomposition and Linear Prediction," submitted to Proceedings of the IEEE.
11. M.P. Ekstrom and J.W. Woods, "2-D Spectral Factorization with Applications in Recursive Digital Filtering," IEEE Trans. Acoust., Speech, and Signal Processing, ASSP-24, 115 (1976).
12. A. Papoulis, Probability, Random Variables, and Stochastic Processes (McGraw-Hill, New York, 1965).
13. K.J. Astrom and P. Eykhoff, "System Identification - A Survey," Automatica 7, 123 (1971).
14. Semiannual Technical Summary, Multi-Dimensional Signal Processing Research Program, Lincoln Laboratory, M.I.T. (31 March 1982), DTIC AD-A118186.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER ESD-TR-82-109	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) Object Detection by Two-Dimensional Linear Prediction		5. TYPE OF REPORT & PERIOD COVERED Technical Report
		6. PERFORMING ORG. REPORT NUMBER Technical Report 632
7. AUTHOR(s) Thomas F. Quatieri		8. CONTRACT OR GRANT NUMBER(s) F19628-80-C-0002
9. PERFORMING ORGANIZATION NAME AND ADDRESS Lincoln Laboratory, M.I.T. P.O. Box 73 Lexington, MA 02173-0073		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS Program Element No. 62702F Project No. 4594
11. CONTROLLING OFFICE NAME AND ADDRESS Air Force Systems Command, USAF Andrews AFB Washington, DC 20331		12. REPORT DATE 28 January 1983
		13. NUMBER OF PAGES 84
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) Electronic Systems Division Hanscom AFB, MA 01731		15. SECURITY CLASS. (of this report) Unclassified
		15a. DECLASSIFICATION DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release; distribution unlimited.		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES None		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) <div style="display: flex; justify-content: space-between;"> <div> object detection image analysis significance testing </div> <div> two-dimensional linear prediction adaptive filtering </div> </div>		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) An important component of any automated image analysis system is the detection and classification of objects. In this report, we consider the first of these problems where the specific goal is to detect anomalous areas (e.g., man-made objects) in textured backgrounds such as trees, grass, and fields of aerial photographs. Our detection algorithm relies on a significance test which adapts itself to the changing background in such a way that a constant false alarm rate is maintained. Furthermore, this test has a potentially practical implementation since it can be expressed in terms of the residuals of an adaptive two-dimensional linear predictor. The algorithm is demonstrated with both synthetic and real-world images.		